

Requested Patent: GB2336921A

Title:

METHOD AND SYSTEM FOR RECOVERING FROM OPERATING SYSTEM
FAILURE ;

Abstracted Patent: US6167494 ;

Publication Date: 2000-12-26 ;

Inventor(s):

HOGGARTH ROGER PHILIP (GB); KNOX RICHARD IAN (GB); CHESTON
RICHARD W (US); LOCKER HOWARD J (US); RHOADES DAVID BENSON (US) ;

Applicant(s): IBM (US) ;

Application Number: US19980067630 19980428 ;

Priority Number(s): US19980067630 19980428 ;

IPC Classification: G06F12/00 ; G06F12/14 ;

Equivalents:

ABSTRACT:

A method, system and computer program product in which half of the storage space on a non-volatile storage device is used to maintain a copy of the last "known-good" copy of the operating system and on invocation by a user, restoring from this backup copy when, for example, the operating system has become corrupted due to changes made to settings or device drivers.

(12) UK Patent Application (19) GB (11) 2 336 921 (13) A

(43) Date of A Publication 03.11.1999

(21) Application No 9821445.5

(22) Date of Filing 05.10.1998

(30) Priority Data

(31) 09/067630 (32) 28.04.1998 (33) US

(71) Applicant(s)

International Business Machines Corporation
(Incorporated in USA - New York)
Armonk, New York 10504, United States of America

(72) Inventor(s)

Richard W Cheston
Roger Philip Hoggarth
Richard Ian Knox
Howard J Locker
David Benson Rhoades

(74) Agent and/or Address for Service

C J Ling
IBM United Kingdom Limited, Intellectual Property
Department, Mail Point 110, Hursley Park,
WINCHESTER, Hampshire, SO21 2JN,
United Kingdom

(51) INT CL⁶

G06F 11/14

(52) UK CL (Edition Q)

G4A AME

(56) Documents Cited

EP 0767431 A1 WO 95/22794 A1
MacUser Vol. 7, No. 4, April 1991, pages 42-55

(58) Field of Search

UK CL (Edition P) G4A AER AME
INT CL⁶ G06F 11/14 11/16
Online: WPI, INSPEC, COMPUTER

(54) Abstract Title

Recovery of data on computer non-volatile storage

(57) Recovery from corruption of data stored on a computer non-volatile storage device, in particular when the operating system has become corrupted due to changes made to settings or device drivers, and the computer is incapable of booting up. The computer non-volatile memory is partitioned into at least two partitions, which may occupy the same hard disk or different disks. The first partition stores the data, a "last known good" copy of which is backed up to the second partition, for example immediately after the system is first installed or immediately after a successful bootup. The backup application program copies the data sector by sector, fig. 3. The POST/BIOS code held in ROM provides a restoration function which, when invoked via a special user key input combination, copies each sector of the second partition back to the corresponding sector of the first partition, fig. 4. Thus the computer is able to boot from this "known good" configuration.




FIG. 3

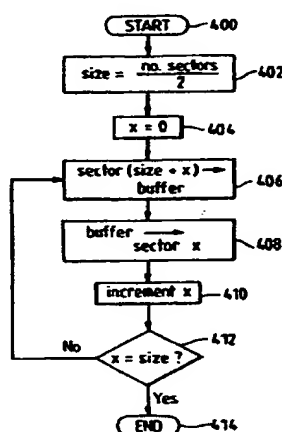


FIG. 4

GB 2 336 921 A







FIG 1 PRIOR ART

FIG. 2FIG. 3

FIG. 4

INTRA-DISK BACKUP

Field of the Invention

5 The present invention relates to recovery of data on computer non-volatile storage, particularly where the corruption of the data causes the operating system of the computer to be corrupted and the computer to be incapable of booting up.

Background of the Invention

10 Personal computer systems are well known in the art. Personal computer systems in general, and IBM Personal Computers in particular, have attained widespread use for providing computer power to many
15 segments of today's modern society. These systems are designed primarily to give independent computer power to a single user and are inexpensively priced for purchase by individuals or small businesses. Personal
20 computers can typically be defined as desktop, floor standing, or portable computers that consist of a system unit having a single central processing unit (CPU) and associated volatile and non-volatile memory, including RAM and BIOS ROM.

25 The capacity of disk drives used in such computers is now so high that it is unlikely that an average end user will ever require all the space available. However, a problem which still remains is that it is very easy for the end user to change some setting or device driver, particularly in the operating system of the computer, which prevents the
30 computer booting successfully. Even if the end user remembers exactly what was changed, it may not be possible to undo the change if the end user has had to boot from a diskette drive to restart the computer. In many computers, booting from a diskette drive provides only a command
35 prompt and from such a command prompt it may not be obvious which files have changed. The problem is potentially much worse for laptop users because their installation disks may be at another location such as at home or in a remote office.

40 European patent EP 0 767 341 discloses a method of backing up a computer disk to another backup medium using the operating system to read a set of logically contiguous sectors from a primary store and write them to a backup medium. The data is restored from the backup medium using a similar set of operating system calls. Mapping is performed by the

operating system to take into account physical flaws on the media. The method in this patent relies on the operating system functioning correctly and does not allow restoration of data in order to overcome operating system corruption due to an end user changing a setting or a device driver.

Research Disclosure n.315, "Mirroring of Data on a Partition Basis", July 1990, discloses a technique for mirroring data on a partition by partition basis. Mirroring can be selected for all of the logical partitions of a disk volume, for none of the logical partitions of a disk volume or for all of the logical partitions of a disk volume. The disk mirroring is done by the filing system and so is dependent on the computer being able to boot up and the operating system being able to start.

IBM Technical Disclosure Bulletin v.39, n.8, "Multimedia Byte for Partial Mirroring", DJ Winarski & TY Winarski, August 1996, discloses the use of a multimedia byte in the header information of files. The multimedia byte depicts the source of the application or file. These may be, for example, CD-ROM, Floppy Disk, 8mm Tape, Downloaded from Lan, Application Output or many others. The byte can be used to control partial mirroring, where only those files which cannot easily be retrieved from another source are copied to another media. In order to use the partial mirroring described in this disclosure, the header bytes of files must be read, which means that it is filing system dependent and the operating system must be functional.

IBM Technical Disclosure Bulletin v.36, n.12, "Automatic Swap of the Primary and Secondary Mirrors in a Disk Mirror System", JL Craft, JM Shieh, December 1993, discloses a system in which sequential mirroring using a Primary and a Secondary partition (or a disk) is used. In normal operation, read or write commands are issued first to the Primary mirror. Write commands are then mirrored onto the Secondary mirror. In the event that data cannot be read from the Primary mirror, an attempt is made to read the data from the Secondary mirror. When the system detects that the Primary mirror may be failing, then it switches the roles of the Primary and Secondary mirrors, so as to increase disk access efficiency. The system cannot detect when the operating system is not functioning due to the end user changing settings or device drivers.

IBM Technical Disclosure Bulletin v.35, n.4b, "User Data Area Protection", JW Blackledge, JF LaPenta Jr., September 1992, discloses a technique which avoids over-writing user data so that no inadvertent destruction of data occurs. Systems which use Initial Microcode Load (IML) have system information stored in, for example, the last 3 Megabytes of the hard disk. When system files are restored from a backup diskette, a check is made of the partition signature bytes. If these are a predetermined signature, the restoration program knows that it is possible that IML data may be overwritten and the user is prompted as to whether this should be done.

It would be advantageous if a method of recovering from operating system corruption which did not rely on the operating system functioning sufficiently to perform restoration could be provided.

Disclosure of the Invention

Accordingly the invention provides a method for recovery from corruption of data stored on a computer non-volatile storage device, the method comprising the steps of: partitioning the non-volatile storage device into a plurality of partitions, including at least a first and a second partition, the first partition storing data, the second partition being capable of storing at least as much data as the first partition; and backing up substantially all of the data stored in the first partition to corresponding locations in the second partition; on corruption of data stored in the first partition: on invocation by a user, restoring substantially all of the stored data from the second partition to corresponding locations in the first partition.

In the event of failure of the operating system due to the end user altering a setting or a device driver, which can be regarded as corruption of data, the end user may invoke a restoration routine at a level below the operating system, that restores the operating system (and other data) from the backup copy stored on the second partition.

In a preferred embodiment, the first and second partitions are stored on the same physical disk. Since computer disk storage capacity is now so large, the use of half of the disk storage as a backup does not impact the end user, who still has sufficient storage capacity in the computer. In an alternative embodiment, the first and second partitions are stored on different physical disks. This has the advantage of

providing a backup unaffected by certain hardware failures of the physical disk carrying the first partition, although this does not provide protection against failure of the physical disk carrying the first partition.

5

Preferably the second partition is not accessible to the end user. This may be achieved by using a special type of disk partition which is 'invisible' to the operating system. This partition cannot be viewed, altered or deleted using the FDISK program or any standard operating system command, such as the FORMAT command.

10

Preferably, the invocation by a user is by means of a pre-determined combination of one or more key depressions on a keyboard. This provides a simple means of invoking the restoration function by means of a combination of keys which a user depresses, in a similar way to the invocation of a reboot function when the combination of Ctrl, Alt and Del are depressed or in a similar manner to the use of F1 during initialisation on some personal computers to enter a setup routine or Ctrl and A to enter a diagnostic routine.

15

20

Further preferably, no operating system is running on the computer at the time when the restoring step is being executed. This is achieved by either initiating the restoring step through the BIOS (Basic Input/Output System) or POST (Power On Self Test) code or by effectively causing the computer to re-boot the operating system, the operating system initialisation code being replaced by the restoration code.

25

The invention also provides a data processing system having data stored on a non-volatile storage, the data processing system comprising: non-volatile storage, divided into a plurality of partitions, including at least a first and a second partition, the first partition storing data, the second partition capable of storing at least as much data as the first partition; means for backing up substantially all of the data stored in the first partition to corresponding locations in the second partition; means for restoring substantially all of the stored data from the second partition to corresponding locations in the first partition; and means for invocation, by a user, of the restoring means so as to restore the operating system.

30

35

40

Further provided by the invention is a computer program product for recovery from corruption of data stored on a computer non-volatile

storage device, the program product comprising: means for partitioning the non-volatile storage into a plurality of partitions, including at least a first and a second partition, the first partition storing at data, the second partition capable of storing at least as much data as the first partition; means for backing up substantially all of the stored data from the first partition to corresponding locations in the second partition; and means, responsive to corruption of the data stored in the first partition and responsive to invocation by an end-user, for restoring substantially all of the data from the second partition to corresponding locations in the first partition.

Brief Description of the Drawings

Embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings, in which:

Figure 1 is a block diagram of a prior art computer system in which the present invention may be used;

Figure 2 is a diagram of the hard disk of Figure 1, showing two partitions on the disk;

Figure 3 is a flow diagram of the processing steps performed during the backup process of the present invention; and

Figure 4 is a flow diagram of the processing steps performed during the recovery process of the present invention.

Detailed Description of the Preferred Embodiment

In Figure 1, a prior art computer 110, comprising a system unit 111, a keyboard 112, a mouse 113 and a display 114 are depicted in block diagram form. The system unit 111 includes a system bus or plurality of system buses 121 to which various components are coupled and by which communication between the various components is accomplished. The microprocessor 122 is connected to the system bus 121 and is supported by read only memory (ROM) 123 and random access memory (RAM) 124 also connected to system bus 121. In many typical computers the microprocessors including the 386, 486 or Pentium microprocessors (Intel and Pentium are trademarks of Intel Corp.). However, other microprocessors including, but not limited to, Motorola's family of

microprocessors such as the 68000, 68020 or the 68030 microprocessors and various Reduced Instruction Set Computer (RISC) microprocessors such as the PowerPC chip manufactured by IBM, or other microprocessors from Hewlett Packard, Sun, Motorola and others may be used in the specific computer.

The ROM 123 contains among other code the Basic Input-Output system (BIOS) which controls basic hardware operations such as the interaction between the CPU and the disk drives and the keyboard. The RAM 124 is the main memory into which the operating system and application programs are loaded. The memory management chip 125 is connected to the system bus 121 and controls direct memory access operations including, passing data between the RAM 124 and hard disk drive 126 and floppy disk drive 127. The CD ROM 132 also coupled to the system 121 is used to store a large amount of data, e.g. a multimedia program or presentation. CD ROM 132 may be an external CD ROM connected through an adapter card or it may be an internal CD ROM having direct connection to the motherboard.

Also connected to this system bus 121 are various I/O controllers: the keyboard controller 128, the mouse controller 129, the video controller 130 and the audio controller 131. As might be expected, the keyboard controller 128 provides the hardware interface for the keyboard 112, the mouse controller 129 provides the hardware interface for mouse 113, the video controller 130 is the hardware interface for the display 114, and the audio controller 131 is the hardware interface for the speakers 115a and 115b. An I/O controller 140 such as a Token Ring adapter card enables communication over a network 146 to other similarly configured data processor systems. These I/O controllers may be located on the motherboard or they may be located on adapter cards which plug into the motherboard, either directly or into a riser card. The adapter cards may communicate with the motherboard using a PCI interface, an ISA or EISA interface or other interfaces.

On the hard disk 126, it is possible to create a special type of disk partition which is 'invisible' to the operating system. Such a partition has been used on personal computers such as the Model 95 from International Business Machines Corporation to contain the power-on self-test (POST) and Basic Input/Output System (BIOS) code. This partition cannot be viewed, altered or deleted using the FDISK program or any standard operating system command, such as the FORMAT command.

Figure 2 shows a hard disk 126, having a partition 201 which occupies the first half of the contents of the hard disk and a partition 202 which occupies the second half of the contents of the hard disk. The contents of the first half can be backed up to the second half. This backup of the first half to the second half is performed by the end user using an application program which copies every sector of the first half of the disk to the corresponding sector of the second half. The application program does not need to concern itself with the partitioning or format of the first half of the disk because it needs perform no interpretation of the data which is read. The physical locations of the two halves on the disk may vary from that shown in figure 2, in which the first half is shown as the outer part of the disk. In a second embodiment, two hard disks may be used, with the contents of a first hard disk being regarded as the first half of the contents of the disk and the contents of a second hard disk being regarded as the second half of the contents of the disk.

The end user should only perform such a backup when he/she knows that the first half of the disk contains 'good' data. Suitable times when the first half of the disk is known to contain 'good' data are immediately after the system is first installed or immediately after a successful bootup. Once the data is copied to the second half of the disk, the end user is unable to alter the copied data on the second half of the disk using anything other than the backup application program so the backup copy of the data will be unaffected by any alterations which the end user makes to the setup of the operating system or device drivers because these changes only affect the active half, that is the first half, of the disk.

Figure 3 shows a flow diagram of the processing performed by the backup program. The backup process starts at step 300. At step 302, the variable *size* is set to be equal to the number of sectors on disk divided by two. At step 304, a sector counter *x* is set to 0. The sector counter counts the number of sectors which have been copied from the first half to the second half of the disk. At step 306, a sector numbered *x* is read and stored into a buffer. At step 308, the contents of the buffer are written to a sector numbered (*size*+*x*) and the write is verified to confirm that the data has been correctly written. At step 310, the variable *x* is incremented and tested at step 312 against the variable *size*. If *x* is equal to *size*, then processing terminates at 314, otherwise processing returns to step 306 and the next sector is read.

The POST/BIOS code held in ROM provides the converse function, invoked via a special key combination during bootup, which copies each sector of the second half of the disk (the backup copy) to the corresponding sector of the first half of the disk (the working copy). Because the copy is performed at the sector level, the ROM code needs no knowledge of the partitioning, formatting, operating system or data content of the sectors which it is copying. (This means that different operating system partitions using different file systems e.g. FAT, HPFS or NTFS are supported without additional function in the ROM). When the restore is complete, the data on the disk drive will be exactly the same as when the last backup was performed and the computer will be able to boot from this 'known good' configuration.

Figure 4 shows a flow diagram of the processing performed by the restore program. The restore process starts at step 400. At step 402, the variable size is set to be equal to the number of sectors on disk divided by two. At step 404, a sector counter x is set to 0. The sector counter counts the number of sectors which have been copied from the first half to the second half of the disk. At step 406, a sector numbered $(size+x)$ is read and stored into a buffer. At step 408, the contents of the buffer are written to a sector numbered x and the write is verified to confirm that the data has been correctly written. At step 410, the variable x is incremented and tested at step 412 against the variable size. If x is equal to size, then processing terminates at 414, otherwise processing returns to step 406 and the next sector is read.

Additional space may be required to be set aside to allow for bad blocks on the disk although modern hard disks are usually capable of doing this under the covers using spare sectors which are maintained for this purpose. In all cases however, the second 'half' of the disk must be at least as big as the first 'half'.

It will be seen that because neither the backup nor the restore is concerned with the data content of each sector, the disk may be partitioned to contain different operating systems, boot manager etc.

It will further be understood that the invention also supports operating system extensions which compress data before writing it to disk. This is because the invention makes no attempt to interpret the data stored in the disk sectors so the compression is completely invisible to the invention, the invention only backing up what is

actually stored on the disk and recovering back what was actually stored on the disk.

CLAIMS

1. A method for recovery from corruption of data stored on a computer non-volatile storage device, the method comprising the steps of:

5 partitioning the non-volatile storage device into a plurality of partitions, including at least a first and a second partition, the first partition storing data, the second partition being capable of storing at least as much data as the first partition;

10 backing up substantially all of the data stored in the first partition to corresponding locations in the second partition;

on invocation by a user, restoring substantially all of the stored data from the second partition to corresponding locations in the first partition.

15 2. A method as claimed in claim 1 wherein the first and second partitions are stored on the same physical disk.

3. A method as claimed in claim 1 wherein the first and second partitions are stored on different physical disks.

20 4. A method as claimed in claim 1 wherein the second partition is not accessible to the end user.

25 5. A method as claimed in claim 1 wherein the data is normally accessed at a filing system level or higher and the step of backing up copies data at a level below the filing system level.

30 6. A method as claimed in claim 1 wherein the data is normally accessed at a filing system level or higher and the step of restoring copies data at a level below the filing system level.

35 7. A method as claimed in claim 1 wherein the invocation by a user is by means of a pre-determined combination of one or more key depressions on a keyboard.

8. A method as claimed in claim 1 wherein no operating system is running on the computer at the time when the restoring step is being executed.

40 9. A data processing system having non-volatile storage, the data processing system comprising:

non-volatile storage, divided into a plurality of partitions, including at least a first and a second partition, the first partition storing data, the second partition capable of storing at least as much data as the first partition;

5 means for backing up substantially all of the data stored in the first partition to corresponding locations in the second partition;

means for restoring substantially all of the stored data from the second partition to corresponding locations in the first partition;

10 means for invocation, by a user, of the means for restoring so as to restore the data.

10. A system as claimed in claim 9 wherein the first and second partitions are stored on the same physical disk.

15 11. A system as claimed in claim 9 wherein the first and second partitions are stored on different physical disks.

12. A system as claimed in claim 9 wherein the second partition is not accessible to the end user.

20 13. A system as claimed in claim 9 wherein the data is normally accessed at a filing system level or higher and the backup means copies data at a level below the filing system level.

25 14. A system as claimed in claim 9 wherein the data is normally accessed at a filing system level or higher and the restoring means copies data at a level below the filing system level.

30 15. A system as claimed in claim 9 wherein the invocation by a user is by means of pre-determined combination of one of more key depressions on a keyboard.

35 16. A system as claimed in claim 9 wherein no operating system is running on the computer at the time when the means for restoring is operating.



Application No: GB 9821445.5
Claims searched: 1-16

Examiner: Melanie Gee
Date of search: 30 October 1998

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.P): G4A (AER, AME)

Int Cl (Ed.6): G06F 11/14, 11/16

Other: Online: WPI, INSPEC, COMPUTER

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X, Y	EP 0767431 A1 (STAC), see col. 3 line 42 - col. 5 line 50.	X: 1-3, 5, 6, 8, 9-11, 13, 14, 16 Y: 4 & 12
Y	WO 95/22794 A1 (APPLE COMPUTER), see whole document.	4 & 12
A	MacUser Vol. 7, No. 4, April 1991, J Taylor et al., "Surviving the crash: hard disk recovery", pages 42-55, and IAC Accession No. 10385573.	

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

THIS PAGE BLANK (USPTO)